

Optimally Combining Classifiers Using Unlabeled Data

Akshay Balsubramani

ABALSUBR@UCSD.EDU

Yoav Freund

YFREUND@UCSD.EDU

Dept. of CSE, University of California San Diego, La Jolla, CA 92093, USA

Abstract

We develop a worst-case analysis of aggregation of classifier ensembles for binary classification. The task of predicting to minimize error is formulated as a game played over a given set of unlabeled data (a transductive setting), where prior label information is encoded as constraints on the game. The minimax solution of this game identifies cases where a weighted combination of the classifiers can perform significantly better than any single classifier.

Keywords: Ensemble aggregation, transductive, minimax

1. Introduction

Suppose that we have a finite set, or ensemble, of binary classifiers $\mathcal{H} = \{h_1, h_2, \dots, h_p\}$, with each h_i mapping data in some space \mathcal{X} to a binary prediction $\{-1, +1\}$. Examples $(x, y) \in \mathcal{X} \times \{-1, +1\}$ are generated i.i.d. according to some fixed but unknown distribution \mathcal{D} , where $y \in \{-1, +1\}$ is the class *label* of the example. We write the expectation with respect to \mathcal{D} or one of its marginal distributions as $\mathbb{E}_{\mathcal{D}}[\cdot]$.

Consider a statistical learning setting, in which we assume access to two types of i.i.d. data: a small set of labeled training examples $S = \{(x'_1, y'_1), \dots, (x'_m, y'_m)\}$ drawn from \mathcal{D} and a much larger set of unlabeled test examples $U = \{x_1, \dots, x_n\}$ drawn i.i.d. according to the marginal distribution over \mathcal{X} induced by \mathcal{D} . A typical use of the labeled set is to find an upper bound on the expected error rate of each of the classifiers in the ensemble. Specifically, we assume a set of lower bounds $\{b_i\}_{i=1}^p > 0$ such that the correlation $\text{corr}(h_i) := \mathbb{E}_{\mathcal{D}}[yh_i(x)]$ satisfies $\text{corr}(h_i) \geq b_i$.

If we ignore the test set, then the best we can do, in the worst case, is to use the classifier with the largest correlation (smallest error). This corresponds to the common practice of *empirical risk minimization* (ERM). However, in many cases we can glean useful information from the distribution of the test set that will allow us to greatly improve over ERM.

We motivate this statement by contrasting two simple prediction scenarios, A and B. In both cases there are $p = 3$ classifiers and $n = 3$ unlabeled test examples. The correlation vector is $\mathbf{b} = (1/3, 1/3, 1/3)$; equivalently, the classifier error rates are 33%. Based on that information, the predictor knows that each classifier makes two correct predictions and one incorrect prediction.

So far, both cases are the same. The difference is in the relations between different predictions on the same example. In case A, each example has two predictions that are the same, and a third that is different. In this case it is apparent that the majority vote over the three classifiers has to be correct on all 3 examples, i.e. we can reduce the error from $\frac{1}{3}$ to 0. In case B, all three predictions are equal for all examples. In other words, the three classification rules are exactly the same on the three examples, so there is no way to improve over any single rule.

These cases show that there is information in the unlabeled test examples that can be used to reduce error – indeed, cases A and B can be distinguished without using any labeled examples. In this paper, we give a complete characterization of the optimal worst-case (minimax) predictions given the correlation vector \mathbf{b} and the unlabeled test examples.

Our development does not consider the feature space \mathcal{X} directly, but instead models the knowledge of the p ensemble predictions on U with a $p \times n$ matrix that we denote by \mathbf{F} . Our focus is on how to use the matrix \mathbf{F} in conjunction with the correlation vector \mathbf{b} , to make minimax optimal predictions on the test examples.

The rest of the paper is organized as follows. In Section 2 we introduce some additional notation. In Section 3 we formalize the above intuition as a zero-sum game between the predictor and an adversary, and solve it, characterizing the minimax strategies for both sides by minimizing a convex objective we call the slack function. This solution is then linked to a statistical learning algorithm in Section 4.

In Section 5, we interpret the slack function and the minimax strategies, providing more toy examples following the one given above to build intuition. In Section 6 we focus on computational issues in running the statistical learning algorithm. After discussing relations to other work in Section 7, we conclude in Section 8.

2. Preliminaries

The main tools we use in this paper are linear programming and uniform convergence. We therefore use a combination of matrix notation and the probabilistic notation given in the introduction. The algorithm is first described in a deterministic context where some inequalities are assumed to hold; probabilistic arguments are used to show that these assumptions are correct with high probability.

The ensemble’s predictions on the unlabeled data are denoted by \mathbf{F} :

$$\mathbf{F} = \begin{pmatrix} h_1(x_1) & h_1(x_2) & \cdots & h_1(x_n) \\ \vdots & \vdots & \ddots & \vdots \\ h_p(x_1) & h_p(x_2) & \cdots & h_p(x_n) \end{pmatrix} \in [-1, 1]^{p \times n} \quad (1)$$

The **true labels** on the test data U are represented by $\mathbf{z} = (z_1; \dots; z_n) \in [-1, 1]^n$.

Note that we allow \mathbf{F} and \mathbf{z} to take any value in the range $[-1, 1]$ rather than just the two endpoints. This relaxation does not change the analysis, because intermediate values can be interpreted as the expected value of randomized predictions. For example, a value of $\frac{1}{2}$ indicates $\{+1 \text{ w.p. } \frac{3}{4}, -1 \text{ w.p. } \frac{1}{4}\}$. This interpretation extends to our definition of the correlation on the test set, $\widehat{\text{corr}}_U(h_i) = \frac{1}{n} \sum_{j=1}^n h_i(x_j) z_j$.¹

The labels \mathbf{z} are hidden from the predictor, but we assume the predictor has knowledge of a **correlation vector** $\mathbf{b} \geq \mathbf{0}^n$ such that $\widehat{\text{corr}}_U(h_i) \geq b_i$ for all $i \in [p]$, i.e. $\frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}$. From our development so far, the correlation vector’s components b_i each correspond to a constraint on the corresponding classifier’s test error $\frac{1}{2}(1 - b_i)$.

The following notation is used throughout the paper: $[a]_+ = \max(0, a)$ and $[a]_- = [-a]_+$, $[n] = \{1, 2, \dots, n\}$, $\mathbf{1}^n = (1; 1; \dots; 1) \in \mathbb{R}^n$, and $\mathbf{0}^n$ similarly. Also, write I_n as the $n \times n$ identity matrix. All vector inequalities are componentwise. The probability simplex in d dimensions is

1. We are slightly abusing the term “correlation” here. Strictly speaking this is just the expected value of the product, without standardizing by mean-centering and rescaling for unit variance. We prefer this to inventing a new term.

denoted by $\Delta^d = \{\sigma \geq \mathbf{0}^d : \sum_{i=1}^d \sigma_i = 1\}$. Finally, we use vector notation for the rows and columns of \mathbf{F} : $\mathbf{h}_i = (h_i(x_1), h_i(x_2), \dots, h_i(x_n))^\top$ and $\mathbf{x}_j = (h_1(x_j), h_2(x_j), \dots, h_p(x_j))^\top$.

3. The Transductive Binary Classification Game

We now describe our prediction problem, and formulate it as a zero-sum game between two players: a predictor and an adversary.

In this game, the predictor is the first player, who plays $\mathbf{g} = (g_1; g_2; \dots; g_n)$, a randomized label $g_i \in [-1, 1]$ for each example $\{\mathbf{x}_i\}_{i=1}^n$. The adversary then plays, setting the labels $\mathbf{z} \in [-1, 1]^n$ under ensemble test error constraints defined by \mathbf{b} . The predictor’s goal is to minimize (and the adversary’s to maximize) the *worst-case expected classification error on the test data* (w.r.t. the randomized labelings \mathbf{z} and \mathbf{g}): $\frac{1}{2} (1 - \frac{1}{n} \mathbf{z}^\top \mathbf{g})$. This is equivalently viewed as maximizing worst-case correlation $\frac{1}{n} \mathbf{z}^\top \mathbf{g}$.

To summarize concretely, we study the following game:

$$V := \max_{\mathbf{g} \in [-1, 1]^n} \min_{\substack{\mathbf{z} \in [-1, 1]^n \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top \mathbf{g} \quad (2)$$

It is important to note that we are only modeling “test-time” prediction, and represent the information gleaned from the labeled data by the parameter \mathbf{b} . Inferring the vector \mathbf{b} from training data is a standard application of Occam’s Razor [Blumer et al. \(1987\)](#), which we provide in Section 4.

The minimax theorem (e.g. [Cesa-Bianchi and Lugosi \(2006\)](#), Theorem 7.1) applies to the game (2), since the constraint sets are convex and compact and the payoff linear. Therefore, it has a minimax equilibrium and associated optimal strategies $\mathbf{g}^*, \mathbf{z}^*$ for the two sides of the game, i.e. $\min_{\mathbf{z}} \frac{1}{n} \mathbf{z}^\top \mathbf{g}^* = V = \max_{\mathbf{g}} \frac{1}{n} \mathbf{z}^{*\top} \mathbf{g}$.

As we will show, both optimal strategies are simple functions of a particular *weighting* over the p hypotheses – a nonnegative p -vector. Define this weighting as follows.

Definition 1 (Slack Function and Optimal Weighting) *Let $\sigma \geq 0^p$ be a weight vector over \mathcal{H} (not necessarily a distribution). The vector of **ensemble predictions** is $\mathbf{F}^\top \sigma = (\mathbf{x}_1^\top \sigma, \dots, \mathbf{x}_n^\top \sigma)$, whose elements’ magnitudes are the **margins**. The **prediction slack function** is*

$$\gamma(\sigma, \mathbf{b}) = \gamma(\sigma) := \frac{1}{n} \sum_{j=1}^n \left[\left| \mathbf{x}_j^\top \sigma \right| - 1 \right]_+ - \mathbf{b}^\top \sigma \quad (3)$$

An **optimal weight vector** σ^* is any minimizer of the slack function: $\sigma^* \in \arg \min_{\sigma \geq 0^p} [\gamma(\sigma)]$.

Our main result uses these to describe the solution of the game (2).

Theorem 2 (Minimax Equilibrium of the Game) *The minimax value of the game (2) is $V = -\gamma(\sigma^*)$. The minimax optimal strategies are defined as follows: for all $i \in [n]$,*

$$g_i^* \doteq g_i(\sigma^*) = \begin{cases} \mathbf{x}_i^\top \sigma^* & |\mathbf{x}_i^\top \sigma^*| < 1 \\ \text{sgn}(\mathbf{x}_i^\top \sigma^*) & \text{otherwise} \end{cases} \quad \text{and} \quad z_i^* = \begin{cases} 0 & |\mathbf{x}_i^\top \sigma^*| < 1 \\ \text{sgn}(\mathbf{x}_i^\top \sigma^*) & |\mathbf{x}_i^\top \sigma^*| > 1 \end{cases} \quad (4)$$




Figure 1: The optimal strategies and slack function as a function of the ensemble prediction $\mathbf{x}^\top \sigma^*$.

The proof of this theorem is a standard application of Lagrange duality and the minimax theorem. The minimax value of the game and the optimal strategy for the predictor \mathbf{g}^* (Lemma 8) are our main objects of study and are completely characterized, and the theorem’s partial description of \mathbf{z}^* (proved in Lemma 9) will suffice for our purposes.²

Theorem 2 illuminates the importance of the optimal weighting σ^* over hypotheses. This weighting $\sigma^* \in \arg \min_{\sigma \geq 0^p} \gamma(\sigma)$ is the solution to a convex optimization problem (Lemma 11), and therefore we can efficiently compute it and \mathbf{g}^* to any desired accuracy. The ensemble prediction (w.r.t. this weighting) on the test set is $\mathbf{F}^\top \sigma^*$, which is the only dependence of the solution on \mathbf{F} .

More specifically, the minimax optimal prediction and label (4) on any test set example \mathbf{x}_j can be expressed as functions of the ensemble prediction $\mathbf{x}_j^\top \sigma^*$ on that test point alone, without considering the others. The \mathbf{F} -dependent part of the slack function also depends separately on each test point’s ensemble prediction. Figure 1 depicts these three functions.

4. Bounding the Correlation Vector

In the analysis presented above we assumed that a correlation vector \mathbf{b} is given, and that each component is guaranteed to be a lower bound on the test correlation of the corresponding hypothesis. In this section, we show how \mathbf{b} can be calculated from a labeled training set.

The algorithm that we use is a natural one using uniform convergence: we compute the empirical correlations for each of the p classifiers, and add a uniform penalty term to guarantee that the \mathbf{b} is a lower bound on the correlation of the test data. For each classifier, we consider three quantities:

- The true correlation: $\text{corr}(h) = \mathbb{E}_{\mathcal{D}} [yh(x)]$
- The correlation on the training set of labeled data: $\widehat{\text{corr}}_S(h) = \frac{1}{m} \sum_{i=1}^m h(x'_i)y'_i$
- The correlation on the test set of unlabeled data: $\widehat{\text{corr}}_U(h) = \frac{1}{n} \sum_{i=1}^n h(x_i)z_i$

2. For completeness, Corollary 10 in the appendices specifies z_i^* when $|\mathbf{x}_i^\top \sigma^*| = 1$.

Using Chernoff bounds, we can show that the training and test correlations are concentrated near the true correlation. Specifically, for each individual classifier h we have the two inequalities

$$\begin{aligned}\Pr(\widehat{\text{corr}}_S(h) > \text{corr}(h) + \epsilon_S) &\leq e^{-2m\epsilon_S^2} \\ \Pr(\widehat{\text{corr}}_U(h) < \text{corr}(h) - \epsilon_U) &\leq e^{-2n\epsilon_U^2}\end{aligned}$$

Let δ denote the probability we allow for failure. If we set $\epsilon_S = \sqrt{\frac{\ln(2p/\delta)}{2m}}$ and $\epsilon_U = \sqrt{\frac{\ln(2p/\delta)}{2n}}$, we are guaranteed that *all* the $2p$ inequalities hold concurrently with probability at least $1 - \delta$.

We thus set the correlation bound to:

$$b_i := \widehat{\text{corr}}_S(h_i) - \epsilon_S - \epsilon_U$$

and have with probability $\geq 1 - \delta$ that \mathbf{b} is a good correlation vector, i.e. $\widehat{\text{corr}}_U(h_i) \geq b_i \quad \forall i \in [p]$.

5. Interpretation and Discussion

Given σ , we partition the examples \mathbf{x} into three subsets, depending on the value of the ensemble prediction: the **hedged set** $H(\sigma) := \{\mathbf{x} : |\mathbf{x}^\top \sigma| < 1\}$, the **clipped set** $C(\sigma) := \{\mathbf{x} : |\mathbf{x}^\top \sigma| > 1\}$, and the **borderline set** $B(\sigma) := \{\mathbf{x} : |\mathbf{x}^\top \sigma| = 1\}$. Using these sets, we now give some intuition regarding the optimal choice of \mathbf{g} and \mathbf{z} given in (4), for some fixed σ .

Consider first examples \mathbf{x}_i in $H(\sigma)$. Here the optimal g_i is to predict with the ensemble prediction $\mathbf{x}_i^\top \sigma$, a number in $(-1, 1)$. Making such an intermediate prediction might seem to be a type of calibration, but this view is misleading. The optimal strategy for the adversary in this case is to set $z_i = 0$, equivalent to predicting ± 1 with probability $1/2$ each. The reason that the learner hedges is because if $g_i < \mathbf{x}_i^\top \sigma$, the adversary would respond with $z_i = 1$ and with $z_i = -1$ if $g_i > \mathbf{x}_i^\top \sigma$. In either case, the loss of the predictor would increase. In other words, our ultimate rationale for hedging is not calibration, but rather “defensive forecasting” in the spirit of [Vovk et al. \(2005\)](#).

Next we consider the clipped set $\mathbf{x}_j \in C(\sigma)$. In this case, the adversary’s optimal strategy is to predict deterministically, and so the learner matches the adversary here. It is interesting to note that with all else held equal, increasing the margin $|\mathbf{x}_j^\top \sigma|$ beyond 1 is suboptimal for the learner. Qualitatively, the reason is that while $\mathbf{x}_j^\top \sigma$ continues to increase, the prediction for the learner is clipped, and so the value for the learner does not increase with the ensemble prediction.

5.1. Subgradient Conditions

For another perspective on the result of Theorem 2, consider the subdifferential set of the slack function γ at an arbitrary weighting σ :

$$\partial\gamma(\sigma) = \left\{ \frac{1}{n} \left(\sum_{\mathbf{x}_j \in C(\sigma)} \mathbf{x}_j \text{sgn}(\mathbf{x}_j^\top \sigma) + \sum_{\mathbf{x}_j \in B(\sigma)} c_j \mathbf{x}_j \text{sgn}(\mathbf{x}_j^\top \sigma) \right) - \mathbf{b}, \quad \forall c_j \in [0, 1] \right\} \quad (5)$$

Note that the hedged set plays no role in $\partial\gamma(\sigma)$. Since the slack function $\gamma(\cdot)$ is convex (Lemma 11), the sub-differential set (5) at any optimal weighting σ^* contains $\mathbf{0}$, i.e.,

$$\exists c_j \in [0, 1] \quad \text{s.t.} \quad n\mathbf{b} - \sum_{j: \mathbf{x}_j^\top \sigma^* > 1} \mathbf{x}_j + \sum_{j: \mathbf{x}_j^\top \sigma^* < -1} \mathbf{x}_j = \sum_{j: |\mathbf{x}_j^\top \sigma^*| = 1} c_j \mathbf{x}_j \text{sgn}(\mathbf{x}_j^\top \sigma^*) \quad (6)$$




Figure 2: An illustration of the optimal $\sigma^* \geq 0^p$. The vector $n\mathbf{b}$ is the difference between the sums of two categories of clipped examples: those with high ensemble prediction ($\mathbf{x}^\top \sigma^* > 1$) and low prediction (< -1). The effect of $B(\sigma^*)$ is neglected for simplicity.

The geometric interpretation of this equation is given in Figure 5.1. The optimal weighting σ^* partitions the examples into five sets: hedged, positive borderline and positive clipped, and negative borderline and negative clipped. Taking the difference between the sum of the positive clipped and the sum of the negative clipped examples gives a vector that is approximately \mathbf{b} . By adding a weighted sum of the borderline examples, \mathbf{b} can be obtained exactly.

5.2. Beating ERM Without Clipping

We now make some brief observations about the minimax solution.

First, note that no σ such that $\|\sigma\|_1 < 1$ can be optimal, because in such a case $-\gamma(\sigma) > -\gamma\left(\frac{\sigma}{\|\sigma\|_1}\right)$; therefore, $\|\sigma^*\|_1 \geq 1$.

Next, suppose we do not know the matrix \mathbf{F} . Then $\|\sigma^*\|_1 = 1$. This can be shown by proving the contrapositive. Assuming the negation $1 < \|\sigma^*\|_1 := a$, there exists a vector $\mathbf{x} \in [-1, 1]^p$ such that $\mathbf{x}^\top \sigma^* = \|\sigma^*\|_1 > 1$. If each of the columns of \mathbf{F} is equal to \mathbf{x} , then by definition of the slack function, $-\gamma\left(\frac{\sigma^*}{a}\right) > -\gamma(\sigma^*)$, so σ^* cannot be optimal.

In other words, if we want to protect ourselves against the worst case \mathbf{F} , then we have to set $\|\sigma\|_1 = 1$ so as to ensure that $C(\sigma)$ is empty. In this case, the slack function simplifies to $\gamma(\sigma) = -\mathbf{b}^\top \sigma$, over the probability simplex. Minimizing this is achieved by setting σ_i to be 1 at $\arg \max_{i \in [p]} b_i$

and zero elsewhere. So as might be expected, in the case that \mathbf{F} is unknown, the optimal strategy is simply to use the classifier with the best error guarantee.

This is true because $C(\sigma^*)$ is empty, and the set of all σ such that $C(\sigma)$ is empty is of wider interest. We dub it the *Zero Box Region*: $\text{ZBR} = \{\sigma : C(\sigma) = \emptyset\}$. Another clean characterization of the ZBR can be made by using a duality argument similar to that used to prove Theorem 2.

Theorem 3 *The best weighting in ZBR satisfies $\max_{\substack{|\mathbf{F}^\top \sigma| \leq \mathbf{1}^n, \\ \sigma \geq 0^p}} \mathbf{b}^\top \sigma = \max_{\mathbf{g} \in [-1, 1]^n} \min_{\frac{1}{n} \mathbf{Fz} \geq \mathbf{b}} \frac{1}{n} \mathbf{z}^\top \mathbf{g}$. In*

particular, the optimal $\sigma^ \in \text{ZBR}$ if and only if the hypercube constraint $\mathbf{z} \in [-1, 1]^n$ is superfluous, i.e. when $V = \min_{\frac{1}{n} \mathbf{Fz} \geq \mathbf{b}} \max_{\mathbf{g} \in [-1, 1]^n} \frac{1}{n} \mathbf{z}^\top \mathbf{g}$.*

The ZBR is where the optimal strategy is always to hedge and never to incorporate any clipping. Consider a situation in which the solution is in ZBR, $\sigma^* = \mathbf{1}^p$, and all of the predictions are binary: $\mathbf{F} \in \{-1, +1\}^{p \times n}$. This is an ideal case for our method; instead of the baseline value $\max_i b_i$ obtained when \mathbf{F} is unknown, we get a superior value of $\sum_i b_i$.

In fact, we referred to such a case in the introduction, and we present a formal version here. Take p to be odd and suppose that $n = p$. Then set \mathbf{F} to be a matrix where each row (classifier) and each column (example) contains $(p+1)/2$ entries equal to $+1$ and $(p-1)/2$ entries equal to -1 .³ Finally choose an arbitrary subset of the columns (to have true label -1), and invert all their entries.

In this setup, all classifiers (rows) have the same error: $\mathbf{b} = \frac{1}{p} \mathbf{1}^p$. The optimal weight vector in this case is $\sigma^* = \mathbf{1}^p$, the solution is in ZBR because $|\mathbf{x}^\top \sigma^*| = 1 \forall \mathbf{x}$, and the minimax value is $V = 1$, which corresponds to zero error. Any single rule has an error of $\frac{1}{2} - \frac{1}{p}$, so using \mathbf{F} with p classifiers has led to a p -fold improvement over random guessing!

Of course, this particular case is extremal in some ways; in order to be in ZBR, there must be many cancellations in $\mathbf{F}^\top \sigma^*$. This echoes the common heuristic belief that, when combining an ensemble of classifiers, we want the classifiers to be “diverse” (e.g. [Kuncheva and Whitaker \(2003\)](#)). The above example in fact has the maximal average disagreement between pairs of classifiers for a fixed p . Similar results hold if \mathbf{F} is constructed using independent random draws.

So our formulation recovers ERM without knowledge of \mathbf{F} , and can recover an (unweighted) majority vote in cases where this provides dramatic performance improvements. The real algorithmic benefit of our unified formulation is in automatically interpolating between these extremes.

To illustrate, suppose \mathbf{F} is given in Table 1, in which there are six classifiers partitioned into two blocs, and six equiprobable test examples. Here, it can be seen that the true labeling must be $+$ on all examples.

	A classifiers			B classifiers		
	h_1	h_2	h_3	h_4	h_5	h_6
x_1	-	+	+	+	+	+
x_2	-	+	+	+	+	+
x_3	+	-	+	+	+	+
x_4	+	-	+	+	+	+
x_5	+	+	-	+	+	+
x_6	+	+	-	-	-	-
\mathbf{b}	1/3	1/3	1/3	2/3	2/3	2/3

Table 1: Example with two classifier blocs.

3. For instance, by setting $F_{ij} = 1$ if $(i+j)$ is even, and -1 otherwise.

In this situation, the best single rule errs on x_6 , as does an (unweighted majority) vote over the six classifiers, and even a vote over just the better-performing “B” classifiers. But a vote over the “A” rules makes no errors, and our algorithm recovers it with a weighting of $\sigma^* = (1; 1; 1; 0; 0; 0)$.

5.3. Approximate Learning

Another consequence of our formulation is that predictions of the form $\mathbf{g}(\sigma)$ are closely related to dual optima and the slack function. Indeed, by definition of $\mathbf{g}(\sigma)$, the slack function value $-\gamma(\sigma) = \mathbf{b}^\top \sigma - \frac{1}{n} \|\mathbf{F}^\top \sigma - \mathbf{g}(\sigma)\|_1 \leq \max_{\sigma' \geq 0^p} [\mathbf{b}^\top \sigma' - \frac{1}{n} \|\mathbf{F}^\top \sigma' - \mathbf{g}(\sigma)\|_1]$, which is simply the dual problem (Lemma 7) of the worst-case correlation suffered by $\mathbf{g}(\sigma)$: $\min_{\substack{\mathbf{z} \in [-1, 1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top [\mathbf{g}(\sigma)]$. We

now state this formally.

Observation 4 *For any weight vector $\sigma \geq 0^p$, the worst-case correlation after playing $\mathbf{g}(\sigma)$ is bounded by*

$$\min_{\substack{\mathbf{z} \in [-1, 1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top [\mathbf{g}(\sigma)] \geq -\gamma(\sigma)$$

Observation 4 shows that convergence guarantees for optimizing the slack function directly imply error guarantees on predictors of the form $\mathbf{g}(\sigma)$, i.e. prediction rules of the form in Fig. 1.

5.4. Independent Label Noise

An interesting variation on the game is to limit the adversary to $z_i \in [-\alpha_i, \alpha_i]^n$ for some $\alpha = (\alpha_1; \dots; \alpha_n) \in [0, 1]^n$. This corresponds to assuming a level $1 - \alpha_i$ of independent label noise on example i : the adversary is not allowed to set the label deterministically, but is forced to flip example i 's label independently with probability $\frac{1}{2}(1 - \alpha_i)$.

Solving the game in this case gives the result (proof in appendices) that if we know some of the ensemble's errors to be through random noise, then we can find a weight vector σ that would give us better performance than without such information.

Proposition 5 (Independent Label Noise)

$$\max_{\mathbf{g} \in [-1, 1]^n} \min_{\substack{-\alpha \leq \mathbf{z} \leq \alpha, \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top \mathbf{g} = \max_{\sigma \geq 0^p} \mathbf{b}^\top \sigma - \frac{1}{n} \sum_{j=1}^n \alpha_j \left[\left| x_j^\top \sigma \right| - 1 \right]_+ > \max_{\sigma \geq 0^p} [-\gamma(\sigma)] = V$$

Our prediction tends to clip – predict with the majority vote – more on examples with more known random noise, because it gains in minimax correlation by doing so. This mimics the Bayes-optimal classifier, which is always a majority vote.

Indeed, this statement's generalization to the asymmetric-noise case can be understood with precisely the same intuition. The sign of the majority vote affects the clipping penalty in the same way:

Proposition 6 (Asymmetric Label Noise) *For some $\mathbf{l}, \mathbf{u} \geq \mathbf{0}^n$,*

$$\max_{\mathbf{g} \in [-1, 1]^n} \min_{\substack{-\mathbf{l} \leq \mathbf{z} \leq \mathbf{u}, \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top \mathbf{g} = \max_{\sigma \geq 0^p} \mathbf{b}^\top \sigma - \frac{1}{n} \sum_{j=1}^n \left(u_j \left[x_j^\top \sigma - 1 \right]_+ + l_j \left[-x_j^\top \sigma - 1 \right]_+ \right) > V$$

6. Computational Issues

The learning algorithm we presented has two steps. The first is efficient and straightforward: \mathbf{b} is calculated by simply averaging over training examples to produce $\widehat{\text{corr}}_S(h_i)$.

So our ability to produce \mathbf{g}^* is dependent on our ability to find the optimal weighting σ^* by minimizing the slack function $\gamma(\sigma)$ over $\sigma \geq 0^p$. Note that typically $p \ll n$, and so it is a great computational benefit in this case that the optimization is in the dual.

We discuss two approaches to minimizing the slack function. The most straightforward approach is to treat the problem as a linear programming problem and use an LP solver. The main problem with this approach is that it requires storing all of the examples in memory. As unlabeled examples are typically much more plentiful than labeled examples, this approach could be practically infeasible without further modification.

A different approach that exploits the structure of the equilibrium uses stochastic gradient descent (SGD). The fact that the slack function is convex guarantees that this approach will converge to the global minimum. The convergence rate might be suboptimal, particularly near the intersections of hyperplanes in the piecewise-linear slack function surface. But the fact that SGD is a constant-memory algorithm is very attractive.

Indeed, the arsenal of stochastic convex optimization methods comes into play theoretically and practically. The slack function is a sum of i.i.d. random variables, and has a natural limiting object $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[[|\mathbf{x}^\top \sigma| - 1]_+ \right] - \mathbf{b}^\top \sigma$ amenable to standard optimization techniques.

7. Related Work

Our duality-based formulation would incorporate constraints far beyond the linear ones we have imposed so far, since all our results hold essentially without change in a general convex analysis context. Possible extensions in this vein include other loss functions as in multiclass and abstaining settings, specialist experts, and more discussed in the next section.

Weighted majority votes are a nontrivial ensemble aggregation method that has received focused theoretical attention for classification. Of particular note is the literature on boosting for forming ensembles, in which the classic work of [Schapire et al. \(1998\)](#) shows general bounds on the error of a weighted majority vote $\epsilon_{\text{WMV}}(\hat{\sigma})$ under any distribution $\hat{\sigma}$, based purely on the distribution of a version of the margin on labeled data.

Our worst-case formulation here gives direct bounds on (expected) test error $\epsilon_{\text{WMV}}(\hat{\sigma})$ as well, since in our transductive setting, these are equivalent to lower bounds on the slack function value by [Observation 4](#). As we have abstracted away the labeled data information into \mathbf{b} , our results depend only on \mathbf{b} and the distribution of margins $|\mathbf{x}^\top \sigma|$ among the unlabeled data. Interestingly, [Amini et al. \(2009\)](#) take a related approach to prove bounds on $\epsilon_{\text{WMV}}(\hat{\sigma})$ in a transductive setting, as a function of the average ensemble error $\mathbf{b}^\top \hat{\sigma}$ and the test data margin distribution; but their budgeting is looser and purely deals with majority votes, in contrast to our \mathbf{g} in a hypercube. The transductive setting has general benefits for averaging-based bounds also ([Blum and Langford \(2003\)](#)).

One class of philosophically related methods to ours uses moments of labeled data in the statistical learning setting to find a minimax optimal classifier; notably among linear separators ([Lanckriet et al. \(2001\)](#)) and conditional label distributions under log loss ([Liu and Ziebart \(2014\)](#)). Our formulation instead uses only one such moment and focuses on unlabeled data, and is thereby more

efficiently able to handle a rich class of dependence structure among classifier predictions, not just low-order moments.

There is also a long tradition of analyzing worst-case binary prediction of online sequences, from which we highlight [Feder et al. \(1992\)](#), which shows universal optimality for bit prediction of a piecewise linear function similar to Fig. 1. The work of [Cesa-Bianchi et al. \(1993\)](#) demonstrated this to result in optimal prediction error in the experts setting as well, and similar results have been shown in related settings ([Vovk \(1990\)](#); [Andoni and Panigrahy \(2013\)](#)).

Our emphasis on the benefit of considering global effects (our transductive setting) even when data are i.i.d. is in the spirit of the idea of shrinkage, well known in statistical literature since at least the James-Stein estimator ([Efron and Morris \(1977\)](#)).

8. Conclusions and Open Problems

In this paper we have given a new method of utilizing unlabeled examples when combining an ensemble of classifiers. We showed that in some cases, the performance of the combined classifiers is guaranteed to be much better than that of any of the individual rules.

We have also shown that the optimal solution is characterized by a convex function we call the slack function. Minimizing this slack function is computationally tractable, and can potentially be solved in a streaming model using stochastic gradient descent. The analysis introduced an ensemble prediction $\mathbf{x}^\top \sigma$ similar to the margin used in support vector machines. Curiously, the goal of the optimization problem is to minimize, rather than maximize, the number of examples with large margin.

Directions we are considering for future research include:

- Is there an algorithm that combines the convergence rate of the linear programming approach with the small memory requirements of SGD?
- In problems with high Bayes error, what is the best way to leverage the generalized algorithm which limits the adversary to a sub-interval of $[-1, 1]$?
- Can the algorithm and its analysis be extended to infinite concept classes, and under what conditions can this be done efficiently?
- Allowing the classifiers to abstain can greatly increase the representational ability of the combination. Is there a systematic way to build and combine such “specialist” classifiers?

Acknowledgements

The authors are grateful to the National Science Foundation for support under grant IIS-1162581.

References

- Massih Amini, Nicolas Usunier, and François Laviolette. A transductive bound for the voted classifier with an application to semi-supervised learning. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 65–72, 2009.
- Alexandr Andoni and Rina Panigrahy. A differential equations approach to optimizing regret trade-offs. *arXiv preprint arXiv:1305.1359*, 2013.
- Avrim Blum and John Langford. Pac-mdl bounds. In *Learning Theory and Kernel Machines*, pages 344–357. Springer, 2003.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Occam’s razor. *Information processing letters*, 24(6):377–380, 1987.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.
- Nicolò Cesa-Bianchi, Yoav Freund, David P Helmbold, David Haussler, Robert E Schapire, and Manfred K Warmuth. How to use expert advice. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 382–391. ACM, 1993.
- B. Efron and C. Morris. Stein’s Paradox in Statistics. *Scientific American*, 236:119–127, 1977.
- Meir Feder, Neri Merhav, and Michael Gutman. Universal prediction of individual sequences. *Information Theory, IEEE Transactions on*, 38(4):1258–1270, 1992.
- Ludmila I Kuncheva and Christopher J Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine learning*, 51(2):181–207, 2003.
- Gert Lanckriet, Laurent E Ghaoui, Chiranjib Bhattacharyya, and Michael I Jordan. Minimax probability machine. In *Advances in neural information processing systems*, pages 801–807, 2001.
- Anqi Liu and Brian Ziebart. Robust classification under sample selection bias. In *Advances in Neural Information Processing Systems*, pages 37–45, 2014.
- Robert E Schapire, Yoav Freund, Peter Bartlett, and Wee Sun Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *Annals of statistics*, pages 1651–1686, 1998.
- Robert J. Vanderbei. *Linear programming: Foundations and extensions*, 1996.
- Vladimir Vovk, Akimichi Takemura, and Glenn Shafer. Defensive forecasting. *AISTATS 2005*, pages 365–372, 2005.
- Volodimir G Vovk. Aggregating strategies. In *Proc. Third Workshop on Computational Learning Theory*, pages 371–383. Morgan Kaufmann, 1990.

Appendix A. Proof of Theorem 2

The core duality argument in the proofs is encapsulated in an independently useful supporting lemma describing the adversary's response to a given \mathbf{g} .

Lemma 7 For any $\mathbf{g} \in [-1, 1]^n$,

$$\min_{\substack{\mathbf{z} \in [-1, 1]^n, \\ \frac{1}{n} \mathbf{F} \mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top \mathbf{g} = \max_{\sigma \geq 0^p} \left[\mathbf{b}^\top \sigma - \frac{1}{n} \left\| \mathbf{F}^\top \sigma - \mathbf{g} \right\|_1 \right]$$

Proof We have

$$\min_{\substack{\mathbf{z} \in [-1, 1]^n, \\ \mathbf{F} \mathbf{z} \geq n \mathbf{b}}} \frac{1}{n} \mathbf{z}^\top \mathbf{g} = \frac{1}{n} \min_{\mathbf{z} \in [-1, 1]^n} \max_{\sigma \geq 0^p} \left[\mathbf{z}^\top \mathbf{g} - \sigma^\top (\mathbf{F} \mathbf{z} - n \mathbf{b}) \right] \quad (7)$$

$$\stackrel{(a)}{=} \frac{1}{n} \max_{\sigma \geq 0^p} \min_{\mathbf{z} \in [-1, 1]^n} \left[\mathbf{z}^\top (\mathbf{g} - \mathbf{F}^\top \sigma) + n \mathbf{b}^\top \sigma \right] \quad (8)$$

$$= \frac{1}{n} \max_{\sigma \geq 0^p} \left[- \left\| \mathbf{g} - \mathbf{F}^\top \sigma \right\|_1 + n \mathbf{b}^\top \sigma \right] = \max_{\sigma \geq 0^p} \left[\mathbf{b}^\top \sigma - \frac{1}{n} \left\| \mathbf{F}^\top \sigma - \mathbf{g} \right\|_1 \right] \quad (9)$$

where (a) is by the minimax theorem. ■

Now \mathbf{g}^* and V can be derived.

Lemma 8 If σ^* is defined as in Theorem 2, then for every $i \in [n]$,

$$g_i^* = \begin{cases} [\mathbf{F}^\top \sigma^*]_i & |[\mathbf{F}^\top \sigma^*]_i| < 1 \\ \text{sgn}([\mathbf{F}^\top \sigma^*]_i) & \text{otherwise} \end{cases}$$

Also, the value of the game (2) is V , as defined in Theorem 2.

Proof [Proof of Lemma 8] From Lemma 7, the primal game (2) is equivalent to

$$\max_{\substack{\mathbf{g} \in [-1, 1]^n, \\ \sigma \geq 0^p}} \left[\mathbf{b}^\top \sigma - \frac{1}{n} \left\| \mathbf{g} - \mathbf{F}^\top \sigma \right\|_1 \right] \quad (10)$$

From (10), it is clear that given a setting of σ , the $\mathbf{g}^*(\sigma)$ that maximizes (10) is also the one that minimizes $\left\| \mathbf{g} - \mathbf{F}^\top \sigma \right\|_1$ under the hypercube constraint:

$$g_i^*(\sigma) = \begin{cases} [\mathbf{F}^\top \sigma]_i, & |[\mathbf{F}^\top \sigma]_i| < 1 \\ \text{sgn}([\mathbf{F}^\top \sigma]_i), & \text{otherwise} \end{cases}$$

The optimum σ here is therefore

$$\begin{aligned} \sigma^* &= \arg \max_{\sigma \geq 0^p} \left[\mathbf{b}^\top \sigma - \frac{1}{n} \left\| \mathbf{g}^*(\sigma) - \mathbf{F}^\top \sigma \right\|_1 \right] \\ &= \arg \max_{\sigma \geq 0^p} \left[\mathbf{b}^\top \sigma - \frac{1}{n} \sum_{j=1}^n \left[|\mathbf{x}_j^\top \sigma| - 1 \right]_+ \right] = \arg \min_{\sigma \geq 0^p} [\gamma(\sigma)] \end{aligned}$$

which finishes the proof. ■

A.1. Derivation of \mathbf{z}^*

We can now derive \mathbf{z}^* from Lagrange complementary slackness conditions.

Lemma 9 *If σ^* is defined as in Theorem 2, then for every $i \in [n]$,*

$$z_i^* = \begin{cases} 0, & |[\mathbf{F}^\top \sigma^*]_i| < 1 \\ \text{sgn}([\mathbf{F}^\top \sigma^*]_i), & |[\mathbf{F}^\top \sigma^*]_i| > 1 \end{cases}$$

Proof We first rewrite the game slightly to make complementary slackness manipulations more transparent.

Define $c = \mathbf{1}^{2n}$, $A = [\mathbf{F}, -\mathbf{F}; -I_{2n}] \in [-1, 1]^{(p+2n) \times 2n}$, and $B = [n\mathbf{b}; -\mathbf{1}^{2n}] \in \mathbb{R}^{p+2n}$. This will allow us to reparametrize the problem in terms of $\zeta = ([z_1]_+, \dots, [z_n]_+, [z_1]_-, \dots, [z_n]_-)^\top$.

Now apply the minimax theorem (Cesa-Bianchi and Lugosi (2006), Theorem 7.1) to (2) to yield the minimax dual game:

$$\min_{\substack{\mathbf{z} \in [-1, 1]^n \\ \frac{1}{n} \mathbf{Fz} \geq \mathbf{b}}} \max_{\mathbf{g} \in [-1, 1]^n} \frac{1}{n} \mathbf{z}^\top \mathbf{g} = \min_{\substack{\mathbf{z} \in [-1, 1]^n \\ \mathbf{Fz} \geq n\mathbf{b}}} \frac{1}{n} \|\mathbf{z}\|_1 \quad (11)$$

With the above definitions, (11) becomes

$$\min_{A\zeta \geq B, \zeta \geq 0} \frac{1}{n} c^\top \zeta \quad (12)$$

This is clearly a linear program (LP); its dual program, equal to it by strong LP duality (since we assume a feasible solution exists; Vanderbei (1996)) is

$$\max_{\lambda \geq 0, c \geq A^\top \lambda} \frac{1}{n} B^\top \lambda \quad (13)$$

Denote the solutions to (12) and (13) as ζ^* and λ^* .

By Lemma 8 and the discussion leading up to it, we already know σ^* , and therefore only need establish the dependence of \mathbf{z}^* on σ^* . Applying LP complementary slackness (Vanderbei (1996), Thm. 5.3) to (12) and (13), we get that for all $j \in [2n]$,

$$[c - A^\top \lambda^*]_j > 0 \implies \zeta_j^* = 0 \quad (14)$$

and

$$\lambda_j^* > 0 \implies [A\zeta^* - B]_j = 0 \quad (15)$$

First we examine (14). The condition $[c - A^\top \lambda^*]_j > 0$ can be rewritten as

$$c_j + [\lambda_-^*]_j - [\mathbf{F}^\top \sigma^*; -\mathbf{F}^\top \sigma^*]_j > 0$$

For any example $i \leq n$, if $|[\mathbf{F}^\top \sigma^*]_i| < 1$, then $[c - A^\top \lambda^*]_i \geq c_i - [\mathbf{F}^\top \sigma^*; -\mathbf{F}^\top \sigma^*]_i > 0$ (since $\lambda_-^* \geq 0$). Similarly, $c_{i+n} - [\mathbf{F}^\top \sigma^*; -\mathbf{F}^\top \sigma^*]_{i+n} > 0$. By (14), this means $\zeta_i^* = \zeta_{i+n}^* = 0$, which implies $z_i^* = 0$ by definition of ζ^* . So we have shown that $|[\mathbf{F}^\top \sigma^*]_i| < 1 \implies z_i^* = 0$.

First we examine (14). The condition $[c - A^\top \lambda^*]_j > 0$ can be rewritten as

$$c_j + [\lambda^*]_j - [\mathbf{F}^\top \sigma^*; -\mathbf{F}^\top \sigma^*]_j > 0$$

For any example $i \leq n$, if $|\mathbf{F}^\top \sigma^*|_i < 1$, then $[c - A^\top \lambda^*]_i \geq c_i - [\mathbf{F}^\top \sigma^*; -\mathbf{F}^\top \sigma^*]_i > 0$ (since $\lambda^*_- \geq 0$). Similarly, $c_{i+n} - [\mathbf{F}^\top \sigma^*; -\mathbf{F}^\top \sigma^*]_{i+n} > 0$. By (14), this means $\zeta_i^* = \zeta_{i+n}^* = 0$, which implies $z_i^* = 0$ by definition of ζ^* . So we have shown that $|\mathbf{F}^\top \sigma^*|_i < 1 \implies z_i^* = 0$.

It remains only to prove that $|\mathbf{F}^\top \sigma^*|_i > 1 \implies z_i^* = \text{sgn}([\mathbf{F}^\top \sigma^*]_i)$ for any $i \in [n]$. We first show this is true for $[\mathbf{F}^\top \sigma^*]_i > 1$. In this case, from the constraints we need $c \geq A^\top \lambda^*$, in particular that

$$0 \leq [c - A^\top \lambda^*]_i = [\lambda^*]_i + (c_i - [\mathbf{F}^\top \sigma^*]_i) \quad (16)$$

By assumption, $c_i - [\mathbf{F}^\top \sigma^*]_i = 1 - [\mathbf{F}^\top \sigma^*]_i < 0$. Combined with (16), this means we must have $[\lambda^*]_i > 0$, i.e. $\lambda_{i+p}^* > 0$. From (15), this means that

$$[A\zeta^* - B]_{i+p} = 0 \iff \zeta_i^* = 1$$

Meanwhile,

$$[c - A^\top \lambda^*]_{i+n} = [\lambda^*]_{i+n} + (c_{i+n} - [-\mathbf{F}^\top \sigma^*]_i) \geq c_{i+n} + [\mathbf{F}^\top \sigma^*]_i > 0$$

so from (14), $\zeta_{i+n}^* = 0$. Since $\zeta_i^* = 1$, this implies $z_i^* = 1 = \text{sgn}([\mathbf{F}^\top \sigma^*]_i)$, as desired.

This concludes the proof for examples i such that $[\mathbf{F}^\top \sigma^*]_i > 1$. The situation when $[\mathbf{F}^\top \sigma^*]_i < -1$ is similar, but the roles of the i^{th} and $(i+n)^{\text{th}}$ coordinates are reversed from (16) onwards in the above proof. \blacksquare

By further inspection of the subgradient conditions described in the body of the paper, one can readily show the following result, which complements Theorem 2.

Corollary 10 For examples j such that $|\mathbf{x}_j^\top \sigma^*| = 1$,

$$z_j^* = c_j \text{sgn}(\mathbf{x}_j^\top \sigma^*)$$

where $c_j \in [0, 1]$ are as defined in (6).

Appendix B. Miscellaneous Proofs

Lemma 11 The function $\gamma(\sigma)$ is convex in σ .

Proof To prove Part 1, note that for each j , the term $\left[|\mathbf{x}_j^\top \sigma| - 1\right]_+$ is convex in σ . Therefore, the average of n terms is convex. As the term $-\mathbf{b}^\top \sigma$ is linear, the whole expression $\gamma(\sigma)$ is convex. (This is a special case of the Lagrangian dual function always being concave in the dual variables.) \blacksquare

Proof [Proof of Theorem 3] Let $A = [\mathbf{F}, -\mathbf{F}] \in \mathbb{R}^{p \times 2n}$ and $c = \mathbf{1}^{2n}$. Then the first assertion follows by LP duality:

$$\begin{aligned} \max_{\substack{|\mathbf{F}^\top \sigma| \leq \mathbf{1}^n, \\ \sigma \geq 0^p}} \mathbf{b}^\top \sigma &= \max_{\substack{A^\top \sigma \leq c, \\ \sigma \geq 0^p}} \mathbf{b}^\top \sigma \stackrel{(a)}{=} \min_{\substack{A\zeta \geq \mathbf{b}, \\ \zeta \geq 0^{2n}}} c^\top \zeta = \min_{\substack{\mathbf{F}(\zeta^\alpha - \zeta^\beta) \geq \mathbf{b}, \\ \zeta^\alpha, \zeta^\beta \geq 0^n}} \left[\|\zeta^\alpha\|_1 + \|\zeta^\beta\|_1 \right] = \min_{\mathbf{F}\mathbf{z} \geq \mathbf{b}} \|\mathbf{z}\|_1 \\ &= \min_{\frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}} \frac{1}{n} \|\mathbf{z}\|_1 = \min_{\frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}} \max_{\mathbf{g} \in [-1, 1]^n} \frac{1}{n} \mathbf{z}^\top \mathbf{g} \stackrel{(b)}{=} \max_{\mathbf{g} \in [-1, 1]^n} \min_{\frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}} \frac{1}{n} \mathbf{z}^\top \mathbf{g} \end{aligned}$$

where (a) is by strong LP duality and (b) uses the minimax theorem.

The second assertion follows because $V = -\gamma(\sigma^*) \stackrel{(c)}{=} \mathbf{b}^\top \sigma^* \stackrel{(d)}{=} \max_{\mathbf{g} \in [-1, 1]^n} \min_{\frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}} \frac{1}{n} \mathbf{z}^\top \mathbf{g}$, where (c) uses the definition of ZBR and (d) is due to the first assertion. \blacksquare

Proof [Proof of Prop. 5] The derivation here closely follows that of Lemma 7, except that (9) now instead becomes

$$\frac{1}{n} \max_{\sigma \geq 0^p} \left[-\alpha \left\| \mathbf{g} - \mathbf{F}^\top \sigma \right\|_1 + n \mathbf{b}^\top \sigma \right]$$

which is equal to the final result. \blacksquare

The proof of Prop. 6 is exactly analogous to that of Prop. 5.